


Computer Networks

CS 552

Badri Nath
Rutgers University
badri@cs.rutgers.edu




About us: Management




- Professor: Badri Nath
 - <http://www.cs.rutgers.edu/~badri>
 - badri@cs.rutgers.edu
 - Office hours: Wednesday 2:4PM
- TA: Le Long
 - longtle@cs.rutgers.edu
- Course info
 - <http://www.cs.rutgers.edu/~badri/552.html>

Course Web Page




- Course schedule
- Reading list
- Lecture notes
- Announcements
- Assignments
- Project ideas
- Exams

Sakai Web page




- <https://sakai>
- Will submit reviews online
- Course announcements
- Written Homeworks

Course Goals




- Understand the basic principles of computer networks, in particular the Internet
- Study new concepts, design principles in network protocols and design
- How to do network research
 - How to determine what is important
 - What are the trends
 - Datacenter, wireless, cloud
 - What are the economics, technology that is driving innovation
 - Cost, performance, hardware

Course Materials




- Research papers
 - Links to pdf on Web page
 - Combination of classic and recent work
 - ~50 papers
 - Optional readings
- Recommended textbooks
 - For students not familiar with networking
 - Peterson & Davie (4th edition)
 - Alternative: Kurose & Ross

Reading papers




- Understand /identify the basic idea
- What kind of a paper?
 - Performance, vision, new direction/protocol paper
- Summarize key idea
- +ve aspects of the paper
 - New, breakthrough, incremental,
- -ve aspects of the paper
 - Readability, Assumptions (valid?), scaling issues (does it scale), implementation (has it been implemented), measurements (problems?)

Books




- **Computer Networks: A Systems Approach**, 4th Ed. (2007), by Larry Peterson and Bruce Davie.
- **Computer Networking: A Top-Down Approach Featuring the Internet**, 5th Ed. (2010), by James F. Kurose and Keith W. Ross.
- **TCP/IP Illustrated, Volume 1: The Protocols** by W. Richard Stevens.
- **Unix Network Programming: Networking APIs: Sockets and XTI (Volume 1)** by W. Richard Stevens.
- **Advanced Programming in the Unix Environment** by W. Richard Stevens, Addison-Wesley, 1993.
- I or 2 recommended

Grading




- 25% Paper summaries/reviews/HWs based on Papers
 - A subset of the papers will be assigned for submitting summary/critique
 - All papers assigned should be read as quizzes/Hws will be based on these papers
- 30% Programming project (two-person)
- 15% Mid term
- 30% Final
- Honor code
 - All submitted work should be yours
 - You are all grad students!!

Class Coverage




- Quick overview of undergraduate networking
- Pre requisite: 352 or equivalent
- Students expected to know
 - Link layer, basic IP routing, TCP,
- Focus on Advanced topics in networking
- Course will deal with:
 - Protocol performance
 - Investigate protocol trade-offs, cost models
 - New Workloads, new technologies

Class etiquette



- Cell phones in off position
- No FB status updates in class
- If you need to surf while in class (I prefer not), do not disturb your neighbors
- Stop me anytime to ask questions
 - Prof may not know the answer!!
- This is a graduate class, student participation in class is important
 - Challenge the class, the prof, and ideas in papers

Lecture Topics




Traditional

- Layering
- Internet architecture
- Routing (IP)
- Transport (TCP)
- Protocols (HTTP, DHCP, DNS)

Recent Topics


- Data centers/Cloud
- Mobility/wireless
- Social Networking
- Enterprise Networks
- Security
- Energy



What is a Network?


- Carrier of information between 2 or more entities
- Some carry objects/people (postal, air, surface transport)
- Most important is the services offered
- User expectation of service
 - Latency, cost, reliability, service interface, others
- We focus on computer networks
- Interconnection may be any medium capable of communicating information:
 - copper wire
 - Lasers (optic fibre)
 - Microwave
 - Cable (coax)
 - wireless
 - satellite link
- Example: Ethernet, Wifi, 3G

Why Networks?




- Availability of Resources
 - Resources become available regardless of the user's physical location (server based, peer2peer)
- Load Sharing/utilization
 - Jobs processed on least crowded machine
 - Resources can be shared
- High Reliability
 - Alternative source of supply (multiple copies)
- Human-to-Human Communication
 - e.g., on-line world


What is Internet Technology?



- What is an internet?
 - Network of networks
- What is *the* Internet?
 - A global internet based on the IP protocol
- To what does “Internet technology” refer?
 - Architecture, services, interfaces, and protocols




Sample Internet Applications




- Electronic mail
- WEB
- File transfer, sharing
- Social networks (FB, linkedin, twitter)
- On-line Shopping
- Search
- Resource distribution (hosting)
- Video conferencing
- Games

Impact of the Net on People




- Access to remote information
 - HW assignments from sakai
 - Stock quotes from financial web site
 - Corporate video, news clips, virtual tours
 - Virtual tours of homes, tourist spots, virtual globetrotting
 - Cloud services
- Person to person and group communication
 - email, collaborative tools (chat groups), Instant messaging, online social networks (FB), twitter, foursquare
- Interactive entertainment
 - youtube, music, video clips

Impact of the Net on Society




- The good
 - Access to information (i-commerce), selling goods and services (e-commerce), incredible productivity tool, unified communication tool
- The bad
 - gossip, too much information, net addicts (FB status updates!)
- The ugly
 - Fraud, pornography, threatening e-mail
- But, it is just a mirror of society

Impact on society




- Net neutrality
- Laws and censorship (SOPA ---)
 - Google vs China
 - Wiki going black -- protest web censorship
- Nations' laws and Internet
 - Regulation
- Content creation, ownership, distribution

Internet Players




- Users, people who use the applications
 - Everyone (mom and pop, kids)
 - get something done (hopefully useful)
- Designers
 - You: protocol design and implementation
 - Scale, performance, cost, incremental deployment
- Service Providers
 - Administrators and ISPs
 - Provider-customer versus peer-to-peer
 - Management, revenue, deployment
- Market/business models for the Internet
 - Consumer to consumer (ebay, match.com, craigslist), Business to consumer (amazon, orbitz, google, netflix, hulu), Business to business (getty, harvest, google), Consumer to business (hot jobs, monster)




Internet Traffic

Consumers Leading Way –
58% CAGR in IP Traffic, 2005E-2011E

Global IP Traffic



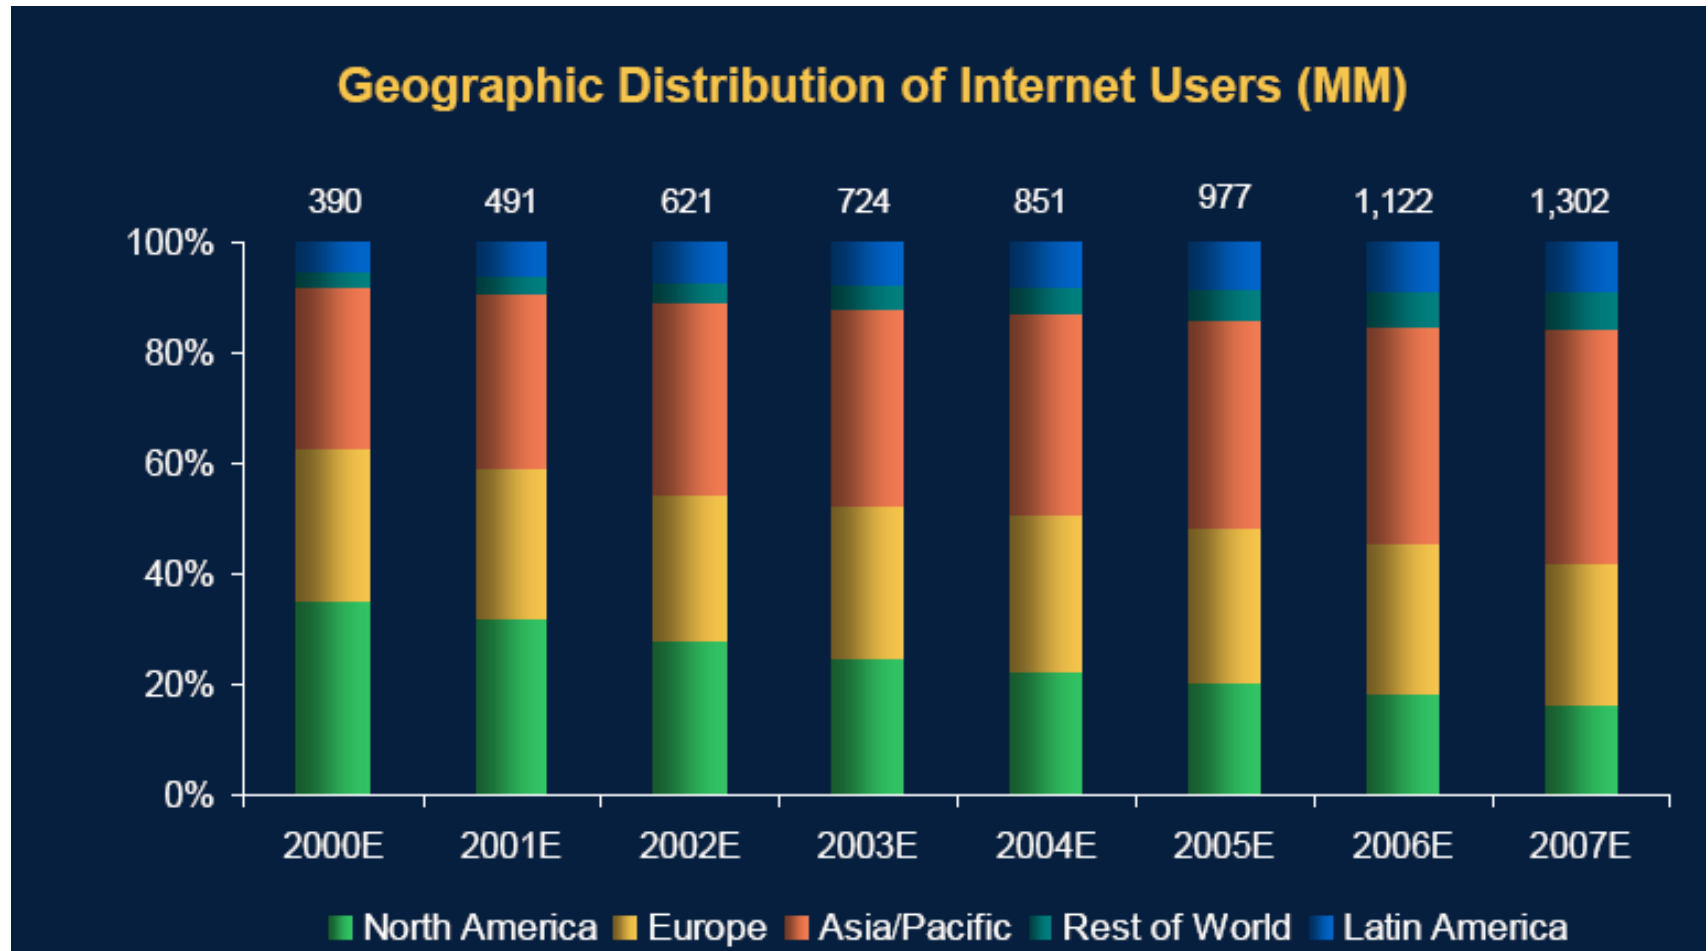

Source: Cisco Systems, Global IP Traffic Forecast and Methodology; Mobility segment (0.1% of traffic in 2007) not displayed





Social Networking





Internet users




Internet transit pricing




WWW Growth

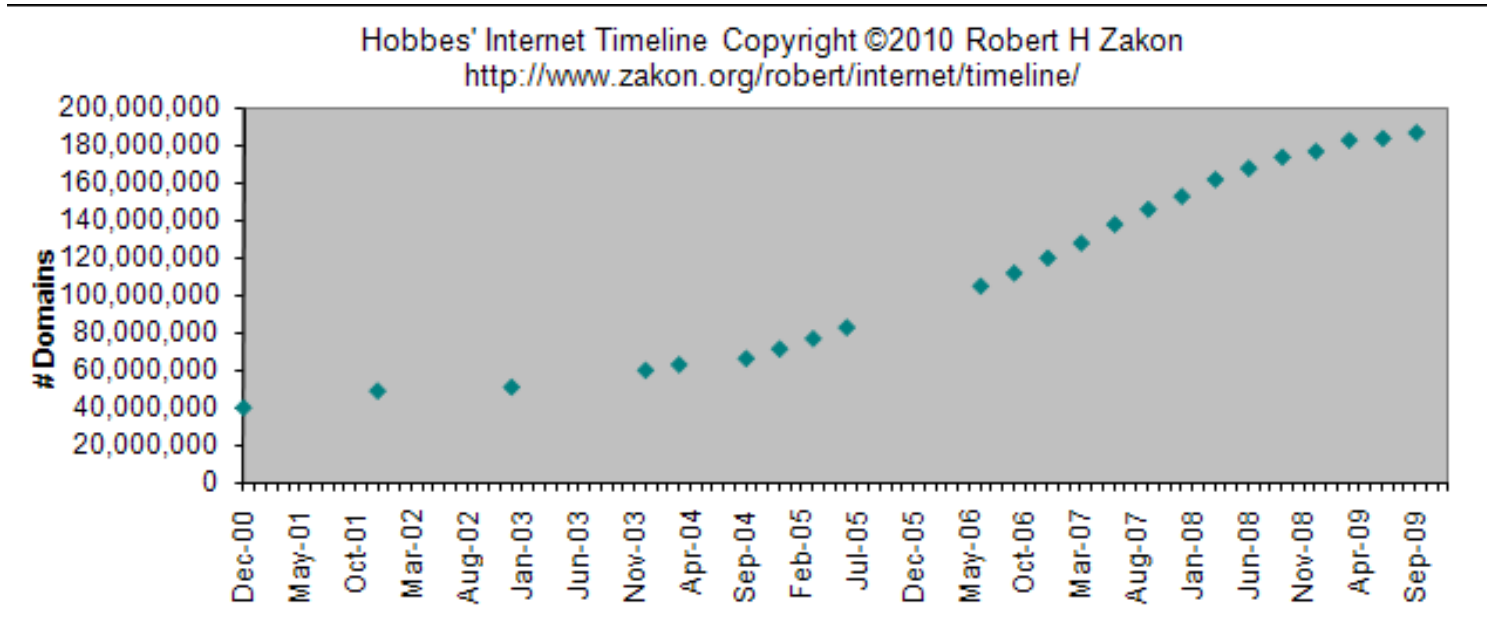





WWW growth




Domain name growth




Facebook growth



Facebook user growth

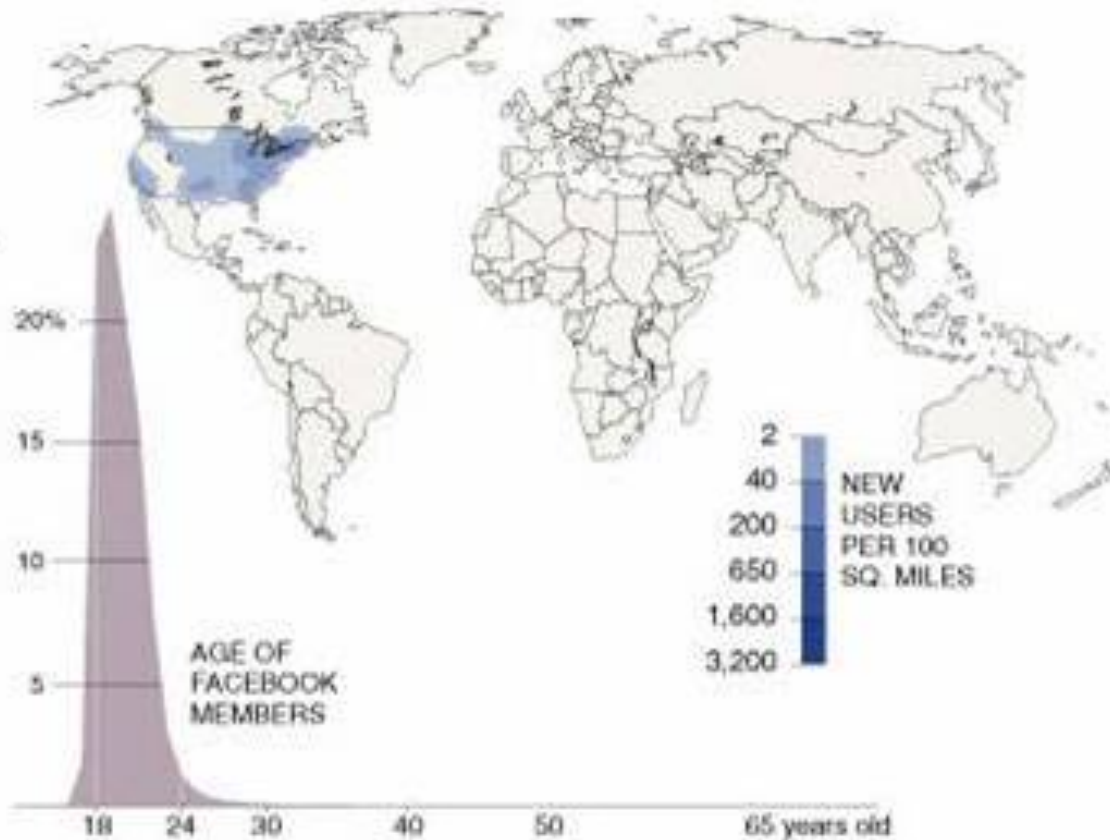




Initially: Only kids!!

February 2004 to
January 2005

Facebook begins at Harvard, and
expands to a few universities at a time.






Now: everyone; grandpa, grandma


February 2008 to
January 2009

Facebook is translated into more than 40 languages. The fastest-growing group of members is people over 35.







Time spent on Internet



Source: Morgan Stanley/comScore Global

Social network usage



Source: Social Networking users & time spent data per comScore, Internet users data per ITU, Facebook data per Facebook.

Copyright 2011. All rights reserved.

Impact of Web 2.0




Alexa Global Traffic Rankings

2005 ⁽¹⁾		2008 ⁽²⁾	
Rank	Web site	Rank	Web site
1	yahoo.com	1	yahoo.com
2	msn.com	2	youtube.com
3	google.com	3	live.com
4	ebay.com	4	google.com
5	amazon.com	5	myspace.com
6	microsoft.com	6	facebook.com
7	myspace.com	7	msn.com
8	google.co.uk	8	hi5.com
9	aol.com	9	wikipedia.org
10	go.com	10	orkut.com


Traffic rank is based on three months of aggregated historical traffic data from Alexa Toolbar users and is a combined measure of page views / users (geometric mean of the two quantities averaged over time).

Today: google, facebook, youtube, twitter, yahoo
Americans conducted 13.9 billion searches in August 2009.
Google Sites accounted for 9 billion searches,
followed by Yahoo! Sites (2.7 billion), Microsoft Sites (1.3 billion),
Ask Network (541 million) and AOL LLC (415 million).





Content Growth



YouTube Content Growth Accelerating (+2x Y/Y) –
35 Hours of Content Added Every Minute



Mobile Internet Usage



Mobile users vs Internet users





Mobile Subscriber base


Country	2006 Mobile Subscriptions (000s)	2006 Growth	Mobile Subscriptions Penetration	Worldwide Share
1 China	445,754	18%	34%	17%
2 USA	232,793	12	78	9
3 Russia	151,937	20	106	6
4 India	137,369	82	12	5
5 Brazil	100,661	17	54	4
6 Japan	94,936	5	74	4
7 Germany	81,242	9	99	3
8 Italy	77,605	11	133	3
9 United Kingdom	69,557	7	115	3
10 Indonesia	58,654	47	26	2
11 Mexico	56,765	22	55	2
12 Turkey	51,659	19	71	2
13 France	51,442	7	84	2
14 Pakistan	48,543	124	31	2
15 Spain	46,339	10	105	2
Total	2,693,087			

Source: Morgan Stanley Research TMT database, Informa




Apple iPhone + iTouch + iPad Ramp – The Likes of Which We Haven't Seen Before

iPhone + iTouch vs. NTT docomo i-mode vs. AOL vs. Netscape Users
First 20 Quarters Since Launch




Morgan Stanley

Note: *AOL subscribers data not available before CQ3:94; Netscape users limited to US only. Morgan Stanley Research estimates ~65MM+ netbooks have shipped in first 11 quarters since launch (10/07). Source: Company Reports, Morgan Stanley Research. Data as of CQ3:10.



Device Growth: connected



Harness Signal from Devices



Site Classification



Search



Connections

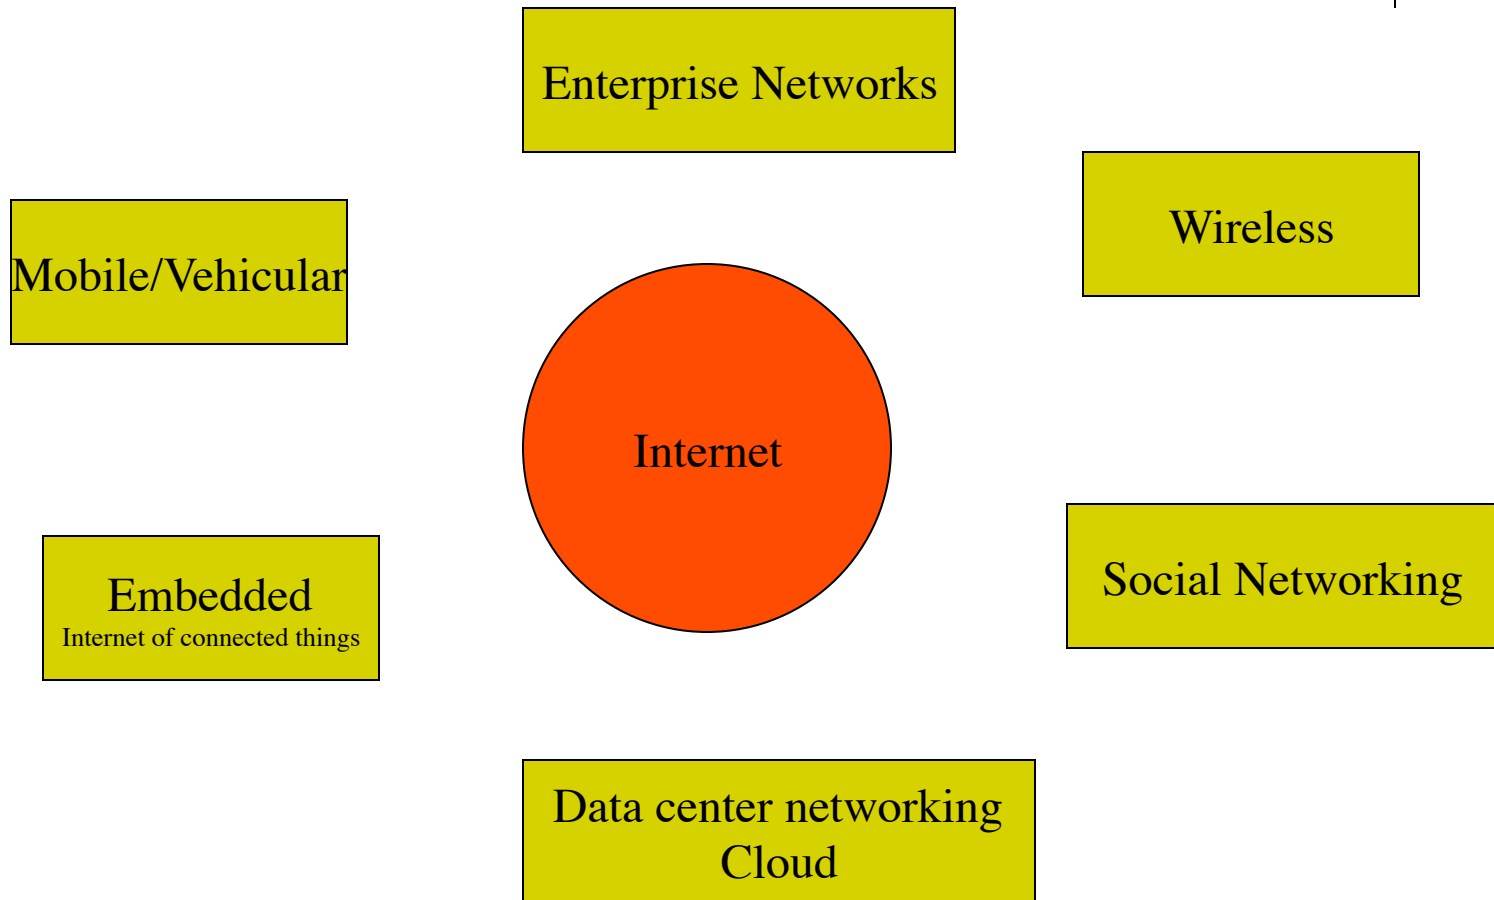



Real-time updates

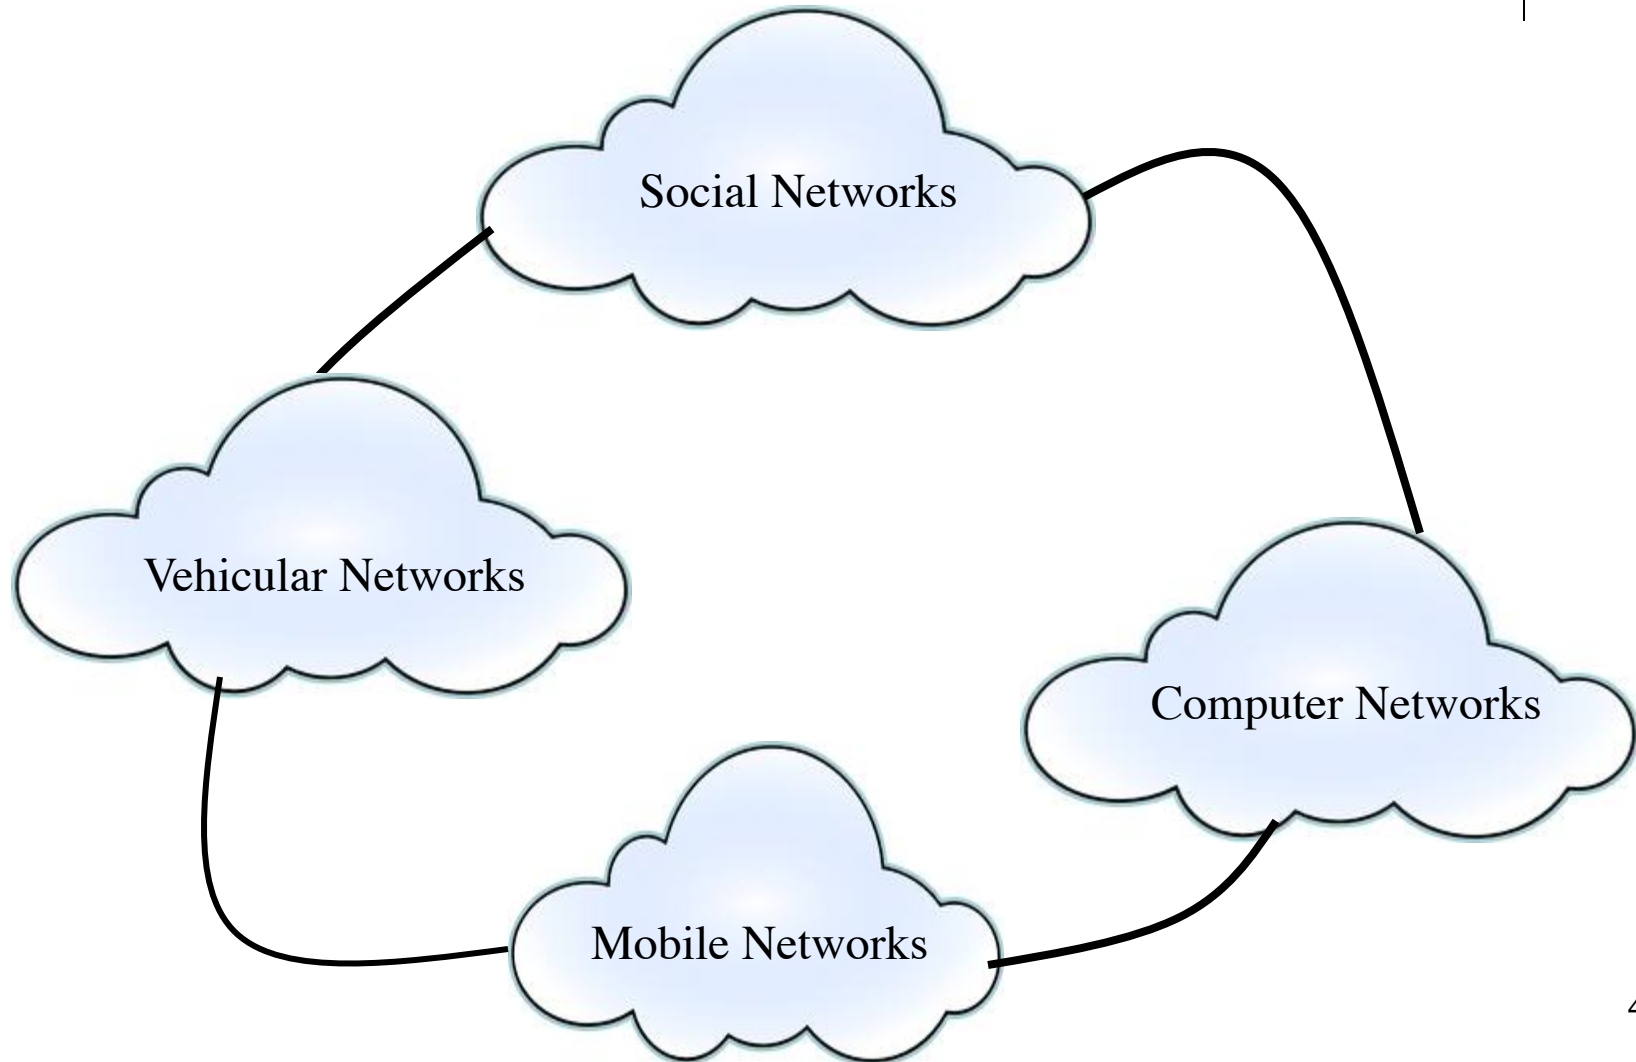

Device State?



What's hot?



Integrating Networks




End point capabilities




- Computer networks
 - Computing nodes
- Social Networks
 - Humans – Crowd sourcing
- Vehicular networks
 - Automobiles
- Mobile networks
 - Smart phones

Social computing




- Can we use humans and computers?
- Exploit humans
 - What can they do best?
- Exploit computers
 - What can they do best?
- Exploit both
 - Socio-computational systems

Objective of networking




- End-hosts to communicate
 - Applications running on end-hosts
- Different technologies
- Different protocols




How to communicate?

- Circuit switching
- Establish a connection before communicating
 - POTS (plain old telephone system)
 - Dedicated pipe for the duration of the session
- Packet switching
 - Multiplex communication from different sources
 - Every packet is self contained
 - Efficient use of resources
 - NO guarantees on performance

How to handle different networks?




- Many differences between networks
- How to translate between various network technologies?
- Have a common protocol for inter network communication
 - IP
 - A set of rules with a well-defined interface



How to locate a node?


- Naming, discovery and routing
 - Network elements needed to support directory
 - Network elements needed to support forwarding towards destination
 - Scalable
 - Reliable

How to meet application demands?



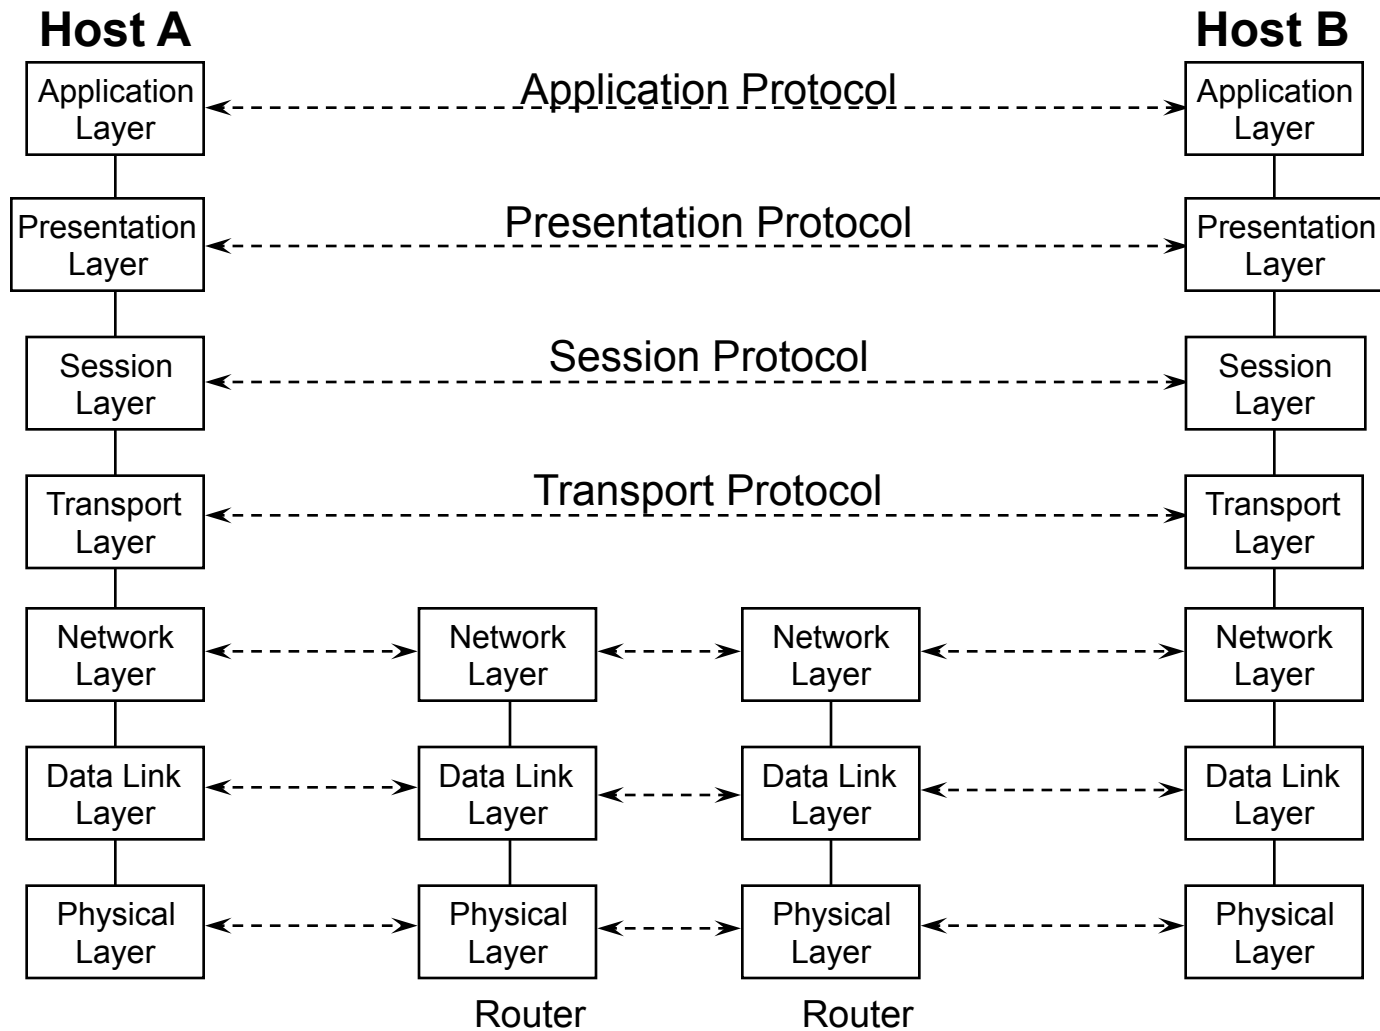

- Corruption?
 - Need error detection and correction
- Reliability
 - Data lost?
- Overload
 - Congestion control
- Security
 - Encryption, authentication

Lots of Functions Needed




- Link
- Multiplexing
- Routing
- Addressing/naming (locating peers)
- Reliability
- Flow control
- Fragmentation
- Etc.....

ISO OSI Layering Architecture





Problems




- Seven layers not widely accepted
- Standardized before implemented
- Top three layers fuzzy
- Internet or TCP/IP layering widespread

TCP/IP Layering Architecture




- A simplified model
- The network layer
 - Hosts drop packets into this layer, layer routes towards destination- only promise- try my best
- The transport layer
 - Reliable/unreliable byte-oriented stream

Internet design philosophy




- #1 Functionality at the edge as opposed to core
- In Telephone network it is the opposite
 - Any new service, the phone company has to provide
 - Edge device is dumb
- Smart device at the edge means programmability
- New services can be supported, drives innovation
 - VOIP (SIP), IM
- Cathedral vs Bazaar

Implications (cathedral vs bazaar)



- Bazaar
 - Edge is programmable
 - Nimble, novel applications
- Cathedral
 - Core elements still rigid
 - Standards, slow evolution
 - Cant do anything radical
- Where is programmability in the system?

2 Multiplexed utilization

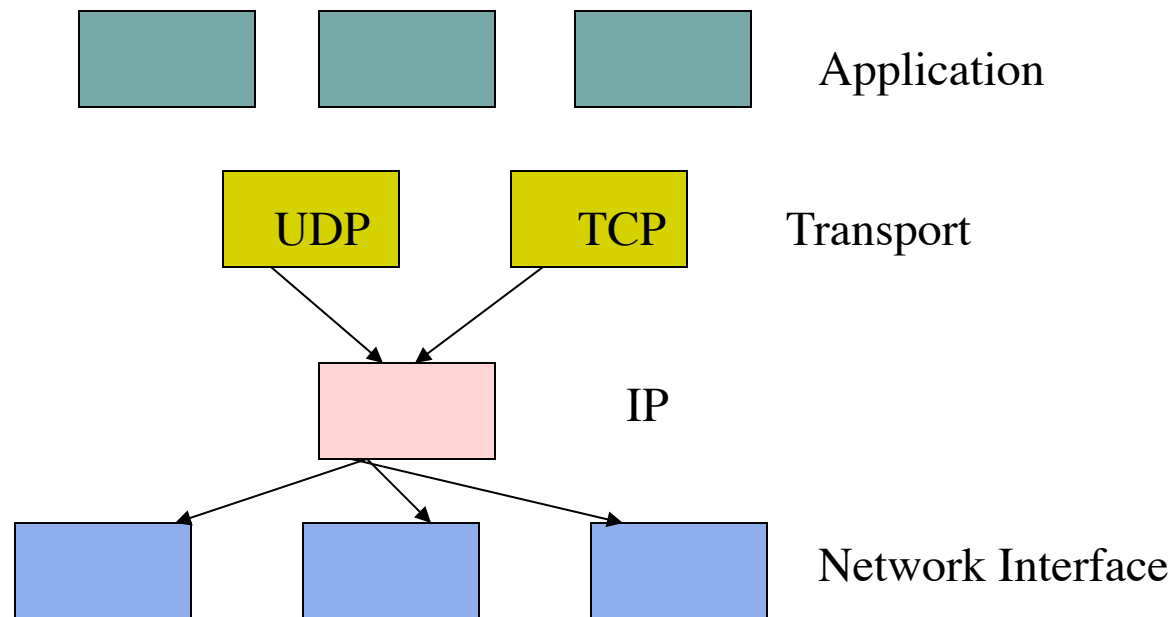


- Best effort, packet switching
- Keep the network simple
 - Packets may be lost, corrupted, out-of-order
 - Let the end host implement any other requirements
 - Want reliability?
 - Retransmit from sender
- Packets self contained
 - Can take different routes
 - Different transfers on the same link
- Stateless in the core
 - End hosts can maintain state
 - Fate-sharing (If I Die, my state will die but not affect others)

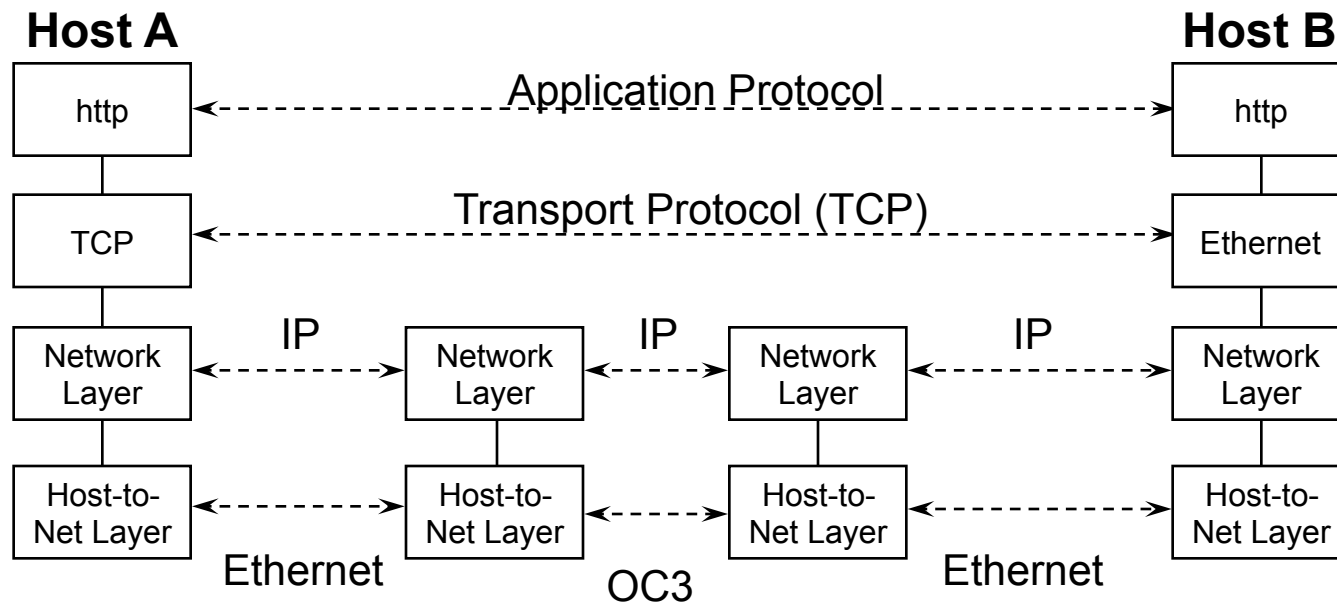

#3 Support multiple networks




- IP over anything, anything over IP
 - Run over any type of link
 - Build any end-to-end protocol over IP



TCP/IP Layering Architecture




#4 distributed management




- Need only know local information
- Information distributed over different nodes
 - Scalable
 - No single hot spot
- Distributed functionality-- Roles
 - Different entities manage different parts of the system
 - Impact on naming, routing, addressing
- Local and Global management authorities

Internet Design Principles




- Scale
 - Protocols should work in networks of all sizes and distances
- Incremental deployment
 - New protocols need to be deployed gradually
- Heterogeneity
 - Different technologies, autonomous organizations
- End-to-end argument
 - Networking functions should be delegated to the edges; application knows best

End-to-end argument



- Saltzer, reed and clark [1984]
 - End-to-end arguments in system design
- Main idea
- If a function can only be completely and correctly implemented with the knowledge and help of the applications standing at the communication end points. Hence providing this function in the subsystem is not possible
 - Complexity at the edges as opposed to the core
- Simply stated, the argument suggests that functions placed at the low levels of a system may be redundant or of little value when compared with the cost of providing them at that low level.
- Don't force feature, service, restriction on the end points

Communication system




- An end system connected by a communication subsystem
- Questions?
- Who is responsible for a given function
 - Subsystem?
 - End units?
 - Or both (redundant) or jointly?

End-to-end argument




- Functions placed at lower level implies specific problems being solved in a general way
- Best aim:
- Simple lower layer with smart end points
 - Basic and general functions at the lower layers
- Gives flexibility

e2e argument




- Low level system may not have all the information to implement the given functionality
- Implement only for performance (wireless links)
- Low level system shared by all applications – what if the application does not need the feature
- Performance
- Duplicated effort
- Should not impact applications that do not use that functionality

e2e tradeoffs




- New business models
 - Network caching, redirection, proxy transcoding
- Wireless Application protocols
 - Gateway provides a box for content translation
- Network redirection
 - Network level switch for load balancing
- Balance between performance, layering, e2e argument

New metrics





- In addition to latency (delay) , bandwidth of a network, other metrics are becoming important
- Energy/power
 - Always-on system consumes a lot of power
- Maintenance
 - Hardware cost is falling, long term human cost of admin is increasing
 - Opex vs capex debate
 - Cloud/Software as a Service (SaaS) models
- Cost of futz
 - Trying to figure what happened
 - Need protocols that reduces burden on end-user
- System Performance/availability
 - How to debug for performance?
 - Where does all the time go? Time scales

Performance tools





- Can we build tools that can answer questions like
 - What is the peak B/W used at a given time scale?
microsec, millisec, sec
 - What is the coarsest time scale where B/W exceeds a threshold L
 - Replace B/W with any other metric CPU usage, DISK B/W, energy usage etc
- *Reference: Efficiently measuring B/W at all time scales, Frank Uyeada et.al, NSDI 2011*

Typical Datacenter networking




Datacenter cost



- 50,000 server @ 3K a pop, 5% cost of money, 3YR
 - 52.5 M/Yr cost
- Power cost
- Power to run the IT equipment
- Power to run cooling, UPS etc – Overhead
- $PUE = \text{Total power} / \text{IT power}$
- 1.2 ideal -- 20% overhead
- Typically 2 to 3 PUE – Air conditioning costs enormous

Facts and Figures [Quereshi09]



- Servers are power hungry (annual electricity bills)


Company	Servers	Electricity	Cost
eBay	16K	$\sim 0.6 \times 10^5$ MWh	$\sim \$3.7$ M
Akamai	40K	$\sim 1.7 \times 10^5$ MWh	$\sim \$10$ M
Rackspace	50K	$\sim 2 \times 10^5$ MWh	$\sim \$12$ M
Microsoft	>200K	$> 6 \times 10^5$ MWh	$> \$36$ M
Google	>500K	$> 6.3 \times 10^5$ MWh	$> \$38$ M
USA (2006)	10.9M	610×10^5 MWh	$\$4.5$ B
MIT campus		2.7×10^5 MWh	$\$62$ M

Energy-proportional metric



- Can we design networks that consume power proportional to utilization?
 - “The Case for Energy-Proportional Computing”, [Luiz André Barroso](#), [Urs Hölzle](#), *IEEE Computer*, vol. 40 (2007).

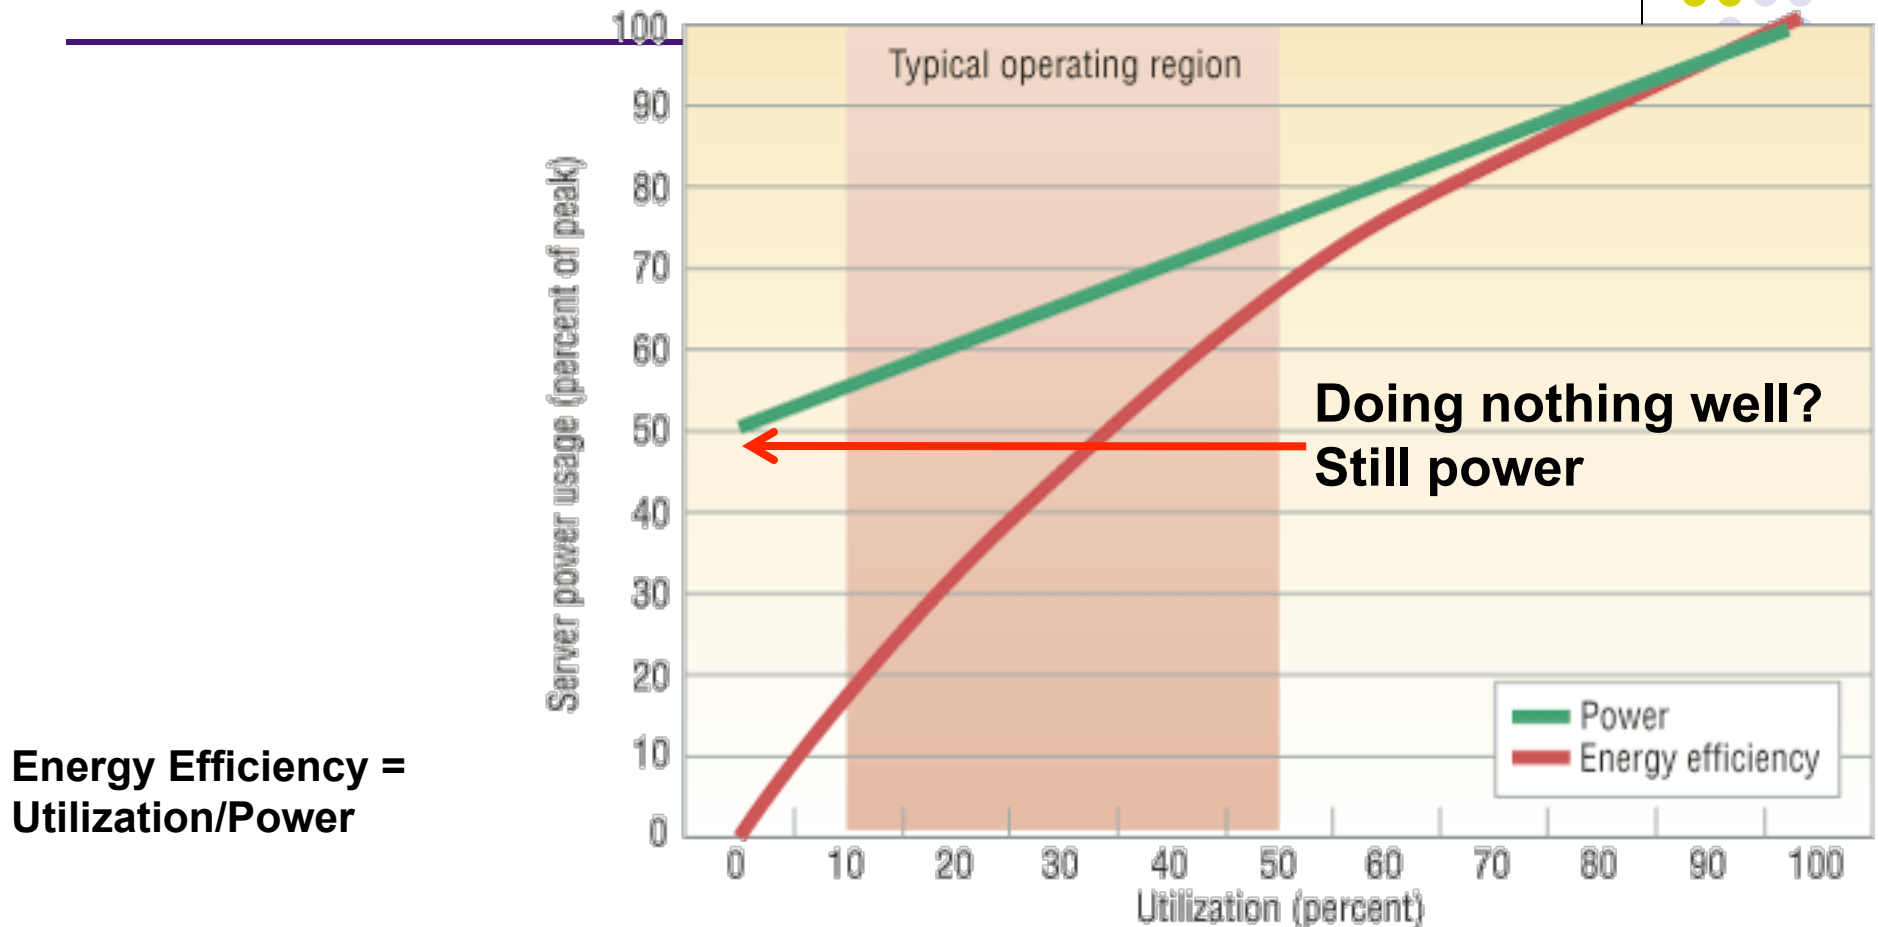
Energy Proportional Computing



“The Case for
Energy-
Proportional
Computing,”
Luiz André
Barroso,
Urs Hölzle,
IEEE Computer
December 2007

Figure 1. Average CPU utilization of more than 5,000 servers during a six-month period. Servers are rarely completely idle and seldom operate near their maximum utilization, instead operating⁷³ most of the time at between 10 and 50 percent of their maximum

Energy Proportional Computing



Energy Efficiency =
Utilization/Power

Figure 2. Server power usage and energy efficiency at varying utilization levels, from idle to peak performance. Even an energy-efficient server still consumes about half its full power when doing virtually no work.

Energy Proportional Computing



Can we do this for
Networking
Infrastructure?

Energy Efficiency =
Utilization/Power





Figure 4. Power usage and energy efficiency in a more energy-proportional server. This server has a power efficiency of more than 80 percent of its peak value for utilizations of 30 percent and above, with efficiency remaining above 50 percent for utilization levels as low as 10 percent.

Energy/ Power metric




- If Cost/Green Conscious
- Make network elements less power hungry
 - [1]Gupta & Singh Greening of the Internet SIGCOMM 2002
 - [2] Energy Proportionality of an enterprise network, Priya Mahadevan, et.al, Green Networking August 2010
- Take advantage of lower power rates elsewhere, time-of-day
 - [3] A. Qureshi, R. Weber, H. Balakrishnan, J. Gutttag, B. Maggs, "Cutting the Electric Bill for Internet-Scale Systems" SIGCOMM 2009

What is the energy consumption of the internet ?



- Energy consumed by networking equipment such as routers, switches, hubs etc
- Does not include hosts
- Internet energy consumption controversial data


Facts and Figures [14]



<i>Device</i>	<i>Approximate Number Deployed</i>	<i>Total AEC TW-h</i>
Hubs	93.5 Million	1.6 TW-h
LAN Switch	95,000	3.2 TW-h
WAN Switch	50,000	0.15TW-h
Router	3,257	1.1 TW-h



$P=E/T$ 1W = 1 J/S 1 Kw-H = 3600000 joules =10 x100W bulbs for 1 Hr

More Numbers...




- Total energy consumed by networking devices annually in 2000 (US): 6.05 TW-h
- Amounts to about 0.07 % of total U.S. energy expenditure
- Expected increase: +1 TW-h by 2005
- Note: This does not include energy consumed by hosts, UPS supplies or cooling equipment.

Enterprise networks



- Network Switches hubs routers under utilized
- Many units not energy proportional
- Turn off ports in proportion to b/w demands
- Timescales?


So, why bother saving energy?



3 reasons:


1. Current energy inefficiencies
 - a) Wired vs. wireless energy costs
 - b) 6 TW-h ~ 1 nuclear reactor
 - c) Extrapolate to World ~140 nuclear reactors
2. Enable greater deployment
 - a) Similar connectivity in India would require 4.75% of total energy budget
3. Enable longer operation times during events of disaster
 - a) Recent Grid failure in NE US/Canada
 - b) Frequent power outages in most of the world

How to save energy ?




- **Low-Energy Hardware Design:**
 - Use hardware components with low power modes of operation
 - Lower the clock frequency of the components, use DVS and other methods during low demand
 - Energy-Aware APIs - give control to software
 - Architecture that allows selective powering off
- **Energy-Aware Protocol Design:**
 - Node-level algorithms for sleeping
 - Route aggregation and other global techniques to inform devices when and for how long to sleep

Where to save energy in a device ?



- Memory
- Main CPU
- Switch fabric or bus backplane
- Line cards (designs ranging from simple to complex with ASICs or network processors to process packets)
- Bang for the buck

Key Questions



- **For how long can components sleep?**
 - state transition times, energy spike
- **How is the decision to sleep taken ?**
 - traffic activity level, in isolation (*uncoordinated*), global (*coordinated*), edge or backbone device, transit or stub network
- **How to distinguish sleep vs. failure ?**
 - should not trigger network reconfiguration in sleep state vs. failure state
- **How to wake up a device ?**
 - at fixed intervals, on packet arrival, account for protocol timers
- **Impact on protocol behavior?**
 - long sleep times, slower propagation of topology changes

More on sleeping...




Uncoordinated sleeping:

- Decision to sleep based on individual traffic levels alone
- Inform nearest neighbors only
- Sleep time limited by protocol hello message timer
- May trigger network reconfiguration in case of missed protocol messages

Coordinated sleeping:

- Decision taken on a network-wide basis
- Need algorithm to pre-compute the optimal sleep time, but computation costs increase
- Hello message frequency can be adjusted, may take longer to detect changes in network topology
- Sleep longer intervals, but forwarding tables may be outdated
- Reroute all traffic through one route, shutdown other routes
- Introduces delay and packet loss in case of sudden traffic burst

Computation Placement



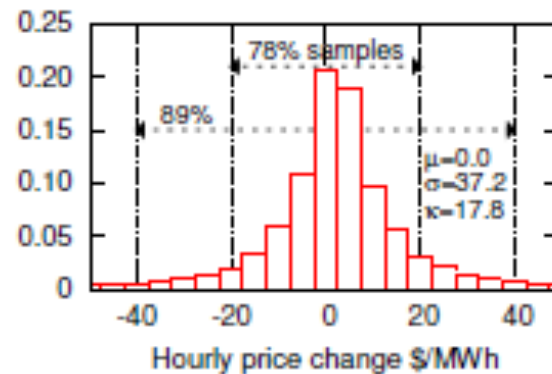
- Energy cost varies by location
- Energy cost varies by hour of the day
- Can we push computation to a geographically distant place to save energy?
- Can we exploit time difference?
- Peak vs non-peak power rates

Exploit spatial-temporal price variation




- Prices vary from place to place
 - Coal vs nuclear
- Price varies with time of day
- Peak vs off-peak

Location	RTO	Mean*	StDev*	Kurt.*
Chicago, IL	PJM	40.6	26.9	4.6
Indianapolis, IN	MISO	44.0	28.3	5.8
Palo Alto, CA	CAISO	54.0	34.2	11.9
Richmond, VA	PJM	57.8	39.2	6.6
Boston, MA	ISONE	66.5	25.8	5.7
New York, NY	NYISO	77.9	40.26	7.9




(a) Palo Alto (NP15)



(b) Chicago (PJM)

Research



- Modification of protocols at layers 2 and 3 to incorporate sleep modes
- Study impact of modifications on end-to-end delay and performance
- Develop energy models for routers and switches
- Study the algorithmic problems of how/when/how long to sleep